

UNITED STATES UTILITY PATENT APPLICATION

FOR

**IMAGE DATA COMPRESSION**

Inventor(s):

Judith A. Goldstein

Lawrence K. Freytag

Michael Keith

Prepared by:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP  
12400 Wilshire Boulevard, 7th Floor  
Los Angeles, California 90025-1026 (408) 720-8598

Express Mail Certificate Of Mailing

"Express Mail" mailing label number EV306655340US Date of Deposit December 16, 2003.  
I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service "Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has been addressed to the Commissioner of Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

## IMAGE DATA COMPRESSION

### BACKGROUND OF THE INVENTION

[0001] This is a divisional of Application Serial No. 09/904,829, filed on July 12, 2001, which is a continuation of Application Serial No. 08/939,321, filed on September 29, 1997, and claims priority thereof.

#### Field of the Invention

[0002] Embodiments of the present invention pertain to the field of image data compression. More particularly, the embodiments pertain to algorithms for near-lossless digital video compression.

#### Background of the Related Art

[0003] A video signal comprises a sequence of frames or images which when displayed at a given frame rate (e.g., 15 to 30 frames per second) simulates the appearance of motion to a human observer. Each frame of the video image comprises a matrix of picture elements, known as "pixels" or "pels." A pixel is the minimum unit of the picture which may be assigned a luminance intensity and a color. In a computer, depending on the data format used, as many as 3 bytes of data can be used to define visual information for a pixel. One popular data format assigns a luminance intensity represented by 1 byte of data to each pixel and further assigns a color represented by 2 bytes of data to groups of 4 pixels each. This format results in an average use of 12 data bits to represent each pixel. Therefore, a single frame at a resolution of 320x240 can be represented by about 120,000 bytes of data.

[0004] Digital video cameras must transfer frame data to a video memory system for display. Multiple frames are transferred over a period of time. The number of frames transferred and

displayed per second is referred to as the frame rate. In general, greater frame rates contribute to heightened appearances of motion, while lower frame rates contribute to the observer being able to perceive individual frames, thus destroying the appearance of motion. A frame rate of 15 frames per second (fps) with a resolution of 320x240 and each pixel being represented by an average of 12 bits requires that about 14 Mbits of data must be transferred between the digital video camera and the video memory each second.

[0005] Digital video cameras are typically connected to a computer system via a peripheral bus. One peripheral bus that is gaining wide acceptance in the computer industry is the Universal Serial Bus (USB). The USB data transfer rate of 8 Mbits/second supports a wide variety of desktop peripherals, from modems, printers, microphones and speakers to graphics tablets, game controls, joysticks, scanners, and digital cameras. However, the USB data rate of 8 Mbits/second is insufficient to support a frame rate of 15 fps for 320x240 video. Further, since several devices may reside on the USB, it is desirable that a digital camera not use the entire 8 Mbits/second.

[0006] Data compression allows an image or video segment to be transferred and stored in substantially fewer bytes of data than required for uncompressed frames. Many methods of digital video compression are based on the idea of eliminating redundant information from frame to frame in a digitized video segment. This is referred to as "interframe compression." Interframe compression methods exploit the temporal redundancy that exists between digital video frames from the same scene recorded moments apart in time. This reduces the required data needed to encode each frame.

[0007] Interframe compression is not ideally suited to the USB environment. This is due to the fact that the USB architecture will not guarantee that every frame of digital video data will be transferred. It is possible that some frames may be dropped. Since interframe compression techniques depend on frame-by-frame redundancies, dropped frames can have a significant impact on picture quality. Therefore, it is desirable to provide a method and apparatus for

compressing digital video images that will allow digital cameras to operate on peripheral busses such as the USB.

[0008] Further, it is desirable to provide a method and apparatus for compressing digital video images that is not perceived as "lossy" by the human visual system. Video compression techniques that humans perceive as "lossy" do not fully preserve all the information originally captured in an image. In general, for "lossy" techniques, as the compression of the video data is increased, the quality of the video suffers.

### BRIEF DESCRIPTION OF THE DRAWINGS

[0009] Figure 1 shows a flow chart of a method for encoding data representing a component of a picture element implemented in accordance with one embodiment of the invention.

[0010] Figure 2 depicts a flow chart of a method for encoding data representing a component of each picture element of a digital image implemented in accordance with one embodiment of the invention.

[0011] Figure 3 shows data representing two scan lines from an example Y-plane of a frame of digital video.

[0012] Figure 4 depicts an example quantization table implemented in accordance with one embodiment of the invention.

[0013] Figure 5a illustrates a high quality, 4-bit encoding example implemented in accordance with one embodiment of the invention.

[0014] Figure 5b depicts an example decoding result from the high quality, 4-bit encoding example illustrated in Figure 5a implemented in accordance with one embodiment of the invention.

[0015] Figure 6a illustrates a high compression, 3-bit encoding example implemented in accordance with one embodiment of the invention.

[0016] Figure 6b depicts an example decoding result from the high compression, 3-bit encoding example illustrated in Figure 6a implemented in accordance with one embodiment of the invention.

[0017] Figure 7 shows an example computer system including an imaging device implemented in accordance with one embodiment of the invention.

[0018] Figure 8a depicts an example 4-bit encode table with no under-correction implemented in accordance with one embodiment of the invention.

[0019] Figure 8b depicts an example 4-bit encode table with 50% under-correction implemented in accordance with one embodiment of the invention.

[0020] Figure 8c depicts an example 4-bit decode table implemented in accordance with one embodiment of the invention.

## DETAILED DESCRIPTION

[0021] A method and apparatus for near-lossless digital video compression is disclosed. In the following description, for the purposes of explanation, specific details are set forth to provide a thorough understanding of the invention. However, it will be obvious to one skilled in the art that these specific details are not required to practice the invention. In other instances, well known methods, devices, and structures are not described in particular detail in order to avoid obscuring the invention.

### Overview

[0022] The invention solves the problem of providing digital image compression that results in an ordinary viewer of the resulting image perceiving the image as "non-lossy" and further that does not rely on interframe redundancies, thereby allowing high quality digital video signals to be transferred across a peripheral bus such as the USB. The invention accomplishes this by using a compression method that takes advantage of redundancies between two scan lines within a given frame of digital video. The compression algorithm of the invention uses Differential Pulse Code Modulation (DPCM) with varying levels or tables of quantizers. In general, and in accordance with one embodiment of the invention, a difference is calculated between a first sample from a current scan line and a corresponding second sample from a previous scan line. A quantization table level is selected from a quantization table that includes at least one level, and information identifying the selected level is placed into an information stream. A quantizer is selected from the quantization table level based on the calculated difference and information corresponding to the selected quantizer is placed into the information stream.

[0023] The invention provides "near-lossless" digital image compression. The term "near-lossless" as used herein means that while some information corresponding to an image is not

preserved during the compression process, an ordinary viewer of the resulting decoded and displayed image would not notice any degradation in quality.

### One Embodiment of the Invention

[0024] For this embodiment, each picture element has a luminance intensity component (the Y component) and two color components (the C<sub>R</sub> and C<sub>B</sub> components). The totality of Y components for a given frame is referred to as the Y plane for that frame. Likewise, the totality of the C<sub>R</sub> and C<sub>B</sub> components for a given frame can be referred to as C<sub>R</sub> and C<sub>B</sub> planes, respectively. These luminance and color components conform to the International Telecommunications Union -Radio Sector (ITU-R) BT.601 standard. Other video color systems, such as RGB, may also be used with the invention.

[0025] Figure 1 shows a flow chart of a method for encoding data representing a component of a picture element implemented in accordance with one embodiment of the invention. At step 110, a difference is calculated between a sample value for a picture element from a current scan line and a sample value for a picture element from a previous scan line.

[0026] Following step 110, a quantization table level is selected at step 120. The quantization table may include any number of levels, with each level including any number of quantizers. An example quantization table is shown in Figure 4, discussed below. The level selection is based on the value of the calculated difference.

[0027] If the selected level differs from a previously selected level, a level switch occurs. This determination is made at step 125. The term "level switch" is more fully discussed below in connection with Figure 2. If there is a level switch, information identifying the selected level is placed into an information stream at step 130. The information stream may be a bit-stream that is transferred across a peripheral bus such as the USB. The information stream may also include

information being transferred across buses with data path widths greater than 1 bit. The information stream is received by a device or system that will perform decoding functions.

[0028] After the information identifying the selected level is placed in the information stream, a quantizer from the selected level is selected at step 140. The term quantizer as used herein is defined as an index into the quantization table. The quantizer selection is based on the value of the calculated difference. Following quantizer selection, information identifying the selected quantizer is placed in the information stream at step 150.

[0029] Figure 2 depicts a flow chart of a method for encoding data representing a component of each picture element of a digital image implemented in accordance with one embodiment of the invention. For this embodiment, each plane is processed independently, although other embodiments are possible where the planes are not processed independently. The method depicted in Figure 2 begins with step 202 where the 0<sup>th</sup> scan line is sampled. The sample values for scan line 0 are stored in a buffer or some other storage device at step 204. Next, the sample values are placed in an information stream at step 206. Note that the 0<sup>th</sup> scan line of each plane is not quantized or encoded at all. The 0<sup>th</sup> line serves as a predictor for the following scan line, as discussed below.

[0030] Following step 206 is step 208 where the 0<sup>th</sup> picture element of the next scan line (line 1) is sampled. After sampling the 0<sup>th</sup> picture element of scan line 1, a difference is calculated at step 210 between the sample value of the 0<sup>th</sup> picture element of line 1 and the sample value of the 0<sup>th</sup> picture element from the 0<sup>th</sup> scan line.

[0031] The present embodiment uses a quantization table that has multiple levels. For example, one level may have quantizers ranging in value from 0 to +/- 7 while another level may have quantizers ranging in value from 0 to +/-181, as shown in the example quantization table of Figure 4 which is discussed below. For the present embodiment, a default level, known to both the encoder and decoder, is used at the beginning of each scan line that is encoded. Since the

default level is known to both the encoder and the decoder, there is no need to communicate the selection of the default level to the decoder. The invention may also be practiced without a default level. In this case, a level must be selected at the beginning of processing each scan line, and the selection of the new level must be communicated to the decoder.

[0032] Following step 210, a quantizer is selected at step 214. The quantizer is selected based on the difference calculated between the sample value of the 0<sup>th</sup> picture element of line 1 and the sample value of the 0<sup>th</sup> picture element from the 0<sup>th</sup> scan line at step 210. The quantizer is selected from the default quantization table level. Following quantizer selection, information identifying the selected quantizer is placed into the information stream at step 216.

[0033] Once the sample value of the 0<sup>th</sup> picture element has been encoded, that is, after the difference has been calculated and after the quantizer has been selected, the quantizer is decoded and written back to a buffer for use as a predictor by the 0<sup>th</sup> element of the next scan line. The decoding occurs at step 218. The decoding step 218 uses the quantizer to look up a quantization value in the currently selected quantization table level. The quantization value is placed in a buffer or other storage device at step 220.

[0034] After processing the 0<sup>th</sup> picture element of line 1 in steps 208, 210, 214, 216, 218, and 220, the next picture element in the current scan line (presently line 1) is processed. The next picture element is sampled at step 222, and a difference between the sample value of the current picture element and the sample value for the corresponding picture element from the previous scan line is calculated at step 224.

[0035] At step 226, a quantization table level is selected. The selection is based on the difference calculated at step 224. Specific example embodiments of level selection algorithms are discussed below in connection with Figure 4.

[0036] Following quantization table level selection at step 226, a determination is made at step 228 on whether a level switch should occur. The term "level switch" as used herein means that a different quantization table level will be used for the current quantizer selection than was used for the previous quantizer selection. The invention may be practiced by allowing level switches every time a level is selected at step 226, or the invention may be practiced by restricting the frequency of level switches. For example, it is possible to only allow level switches every four times a picture element is processed. Thus, a single level is used for at least four contiguous samples before a level switch can occur. Restricting the frequency of level switches has the benefits of reducing traffic on the information stream, helping to provide an adequate compression ratio, and allowing easier and speedier decoding. Further, it is possible to limit the total number of times a level shift can occur for an entire frame. The Y, C<sub>R</sub> and C<sub>B</sub> planes can be considered separately or in combination when determining a level shift maximum. As an example, the maximum number of level shifts can be set to equal 10% of the byte size of one raw (uncompressed) video frame, plus one. The plus one is there in order to handle the case where the 10% limit is reached part way through a scan line. The one extra level switch can be used to switch to the default level, and the default level will then be used for the remainder of the frame. Level shift maximums other than 10% are also possible. Further, the level shift maximum may be varied at any time.

[0037] If a determination is made at step 228 that a level switch is required, information identifying the new level is placed into the information stream at step 230. Step 230 may include placing an escape code into the information stream to let the decoding device know that information identifying a new level follows. Following step 230 is step 232.

[0038] If no level switch is required or allowed, control passes to step 232 following step 228. At step 232, a quantizer is selected from the proper quantization table level and information identifying the selected quantizer is placed into the information stream. Following step 232, the

current sample is decoded at step 234 and the decoded current sample is placed in a buffer or other storage device at step 236.

[0039] If previous steps 222 through 236 processed the last picture element in the current scan line, then control passes to step 240. Otherwise, step 222 follows step 238 and another picture element is processed at steps 222 through 236. These steps are repeated until the final picture element for the current scan line has been processed.

[0040] At step 240, a determination is made on whether the current scan line is the last scan line in the frame. If the current scan line is the last scan line, then the processing of the frame (or at least one plane of the frame) has completed. Otherwise, the processing of the next scan line begins at step 208. Steps 208 through 240 are repeated until the last scan line has been processed.

[0041] Figures 3 through 6b set forth a couple of examples of how the invention may be practiced. Figure 3 shows an example Y-plane 300 consisting of two scan lines and eight sampled picture elements per scan line. The Y -plane 300 is not intended to represent a realistic sampled Y-plane, but is disclosed in order to provide a simple example. The Y-plane 300 is discussed below in connection with Figures 5a through 6b.

[0042] Figure 4 depicts an example quantization table 400 implemented in accordance with one embodiment of the invention. The table 400 consists of 5 levels with levels 0 through 2 including 15 quantizers each and levels 3 and 4 including 7 quantizers each. One embodiment of the invention has two modes of operation: 1) a high quality, 4-bit encoding mode; and 2) a high compression, 3-bit encoding mode. For table 400, the quantizers delineated by the dashed line are used for the high quality mode, while the quantizers delineated by the solid line are used for the high compression mode. The quantization values included in the table 400 are examples. It is possible to practice the invention using different numbers of levels, embodying different numbers of quantizers and using different quantizer values.

[0043] The invention may also be practiced with more than one quantization table. It is possible to switch between two or more quantization tables as applications require. For example, one table could be used to provide standard encoding while another is used to provide under correction. Under correction has the effect of both smoothing out quantization noise (the decoded and displayed images look subtly blurred, but this tends to look "better"), and reducing the number of level shifts. One method of applying under correction involves multiplying the differences calculated in steps 210 and 224 in Figure 2 by a factor of, for example, between .5 and 1. Another method of applying under correction is to scale the quantization table values by a factor of, for example, between 1 and 2. A factor of 1 in each method if applying under correction is the same as having no under correction. Thus, one table could be used to provide no under correction while another table could be used to provide under correction. Multiple tables could be used to provide varying degrees of under correction. Further, while one table is being used for encoding, another table may be updated or its values changed.

[0044] One embodiment of the invention also provides a separate decode table. This allows an under correction factor to be applied to the encode table(s) only. Alternatively, the decode table could include the same quantization values as the encode table(s).

[0045] Figure 5a illustrates a high quality, 4-bit encoding example implemented in accordance with one embodiment of the invention. This example uses the example Y-plane 300 shown in Figure 3 and the quantization table 400 shown in Figure 4. First, the difference between the 0<sup>th</sup> sample of the current line and the 0<sup>th</sup> sample of the previous line is calculated. The difference has a value of -1. Next, a "best" level is selected. The method for selecting the "best" level for this example embodiment is as follows:

For all Y-planes data

if (absolute value of the calculated difference) < 65

    best level = 1; (default level at the beginning of each scan line )

```

        else
            best level = 2;

For all CR and CB plane data

    if (absolute value of calculated difference) < 8
        best level = 0; ( default level at the beginning of each scan line )
    else
        best level = 1.

```

Note that level 0 is not used when encoding the Y-plane and level 2 is not used when encoding the C<sub>R</sub> and C<sub>B</sub> planes. Also, the present embodiment implements a policy of restricting level shifts to once every four samples. One possibility, implemented in the present embodiment, is to determine the "best" level for four contiguous samples, then use the highest level for all four samples. As mentioned above, the default level for the beginning of each scan line for the Y-plane is level 1. Thus, the first four samples in the present example use level I.

[0046] Once the "best" level is determined, the difference (-1 in the case of sample 0) is matched to the closest quantization table 400 value. The closest quantization table 400 value for in level 1 for sample 0 is 0. Quantization value 0 corresponds to quantizer 0, thus 0 is output into the information stream. Samples 1 through 3 are likewise processed, with level 1 used in each case.

[0047] For sample 4, the difference between the 4<sup>th</sup> sample of the previous scan line and the 4<sup>th</sup> sample of the current line is -119. Using the "best" level algorithm disclosed above, the "best" level for the 4<sup>th</sup> sample is 2. Note that the highest "best" level for sample 4 through 7 is level 2. Thus, level 2 will be used for samples 4 through 7. In the present example, an escape code of 15 is placed in the information stream to signal to the decoding device that a new level follows. Thus, a value of 15 followed by a value of 2 are placed into the information stream. The difference of -119 for the 4<sup>th</sup> sample matches most closely in level 2 with the quantization value

[0048] -129 which corresponds with quantizer 12. Therefore, a value of 12 is placed into the information stream. Samples 5 through 7 are processed in similar fashion, each using level 2.

[0049] Figure 5b depicts an example decoding result from the high quality, 4-bit encoding example illustrated in Figure 5a implemented in accordance with one embodiment of the invention. The values in Figure 5b are the result of decoding the information placed into the information stream in the example shown in Figure 5a. The decoding table used to determine the values in Figure 5b includes the same quantization table values as the table 400 in Figure 4.

[0050] Figure 6a illustrates a high compression, 3-bit encoding example implemented in accordance with one embodiment of the invention. This example also uses the example Y-plane 300 shown in Figure 3 and the quantization table 400 shown in Figure 4. First, the difference between the 0<sup>th</sup> sample of the current line and the 0<sup>th</sup> sample of the previous line is calculated. The difference has a value of -1. Next, a "best" level is selected. The method for selecting the "best" level for this example embodiment is as follows:

For all Y-plane data

```
if (absolute value of the calculated difference) < 33  
    best level = 2; (default level at the beginning of each scan line)  
else if (absolute value of the calculated difference) < 65  
    best level = 3 ;  
else  
    best level = 4;
```

For all C<sub>R</sub> and C<sub>B</sub> plane data

```
if (absolute value of calculated difference) < 11  
    best level = 0; ( default level at the beginning of each scan line )  
else if (absolute value of calculated difference) < 21  
    best level = 1 ;  
else  
    best level = 2.
```

Note that levels 0 and 1 are not used when encoding the Y-plane and levels 3 and 4 are not used when encoding the C<sub>B</sub> and C<sub>R</sub> planes. This example also implements the policy of restricting level shifts to once every four samples, and in particular, the policy of using the highest "best" level for four contiguous samples. As mentioned above, the default level for the beginning of each scan line for the Y-plane is level 2. Thus, the first four samples in the present example use level 2.

[0051] Once the "best" level is determined, the difference (-1 in the case of sample 0) is matched to the closest quantization table 400 value. The closest quantization table 400 value for in level 2 for sample 0 is 0. Quantization value 0 corresponds to quantizer 0, thus 0 is output into the information stream. Samples 1 through 3 are likewise processed, with level 2 used in each case.

[0052] For sample 4, the difference between the 4<sup>th</sup> sample of the previous scan line and the 4<sup>th</sup> sample of the current line is -119. Using the "best" level algorithm disclosed above, the "best" level for the 4<sup>th</sup> sample is 4. Note that the highest "best" level for sample 4 through 7 is level 4. Thus, level 4 will be used for samples 4 through 7. In the present example, an escape code of 7 is placed in the information stream to signal to the decoding device that a new level follows. Thus, a value 7 followed by a value 4 are placed into the information stream. The difference of -119 for the 4<sup>th</sup> sample matches most closely in level 4 with the quantization value -128 which corresponds with quantizer 6. Therefore, a value 6 is placed into the information stream. Samples 5 through 7 are processed in similar fashion, each using level 4.

[0053] Figure 6b depicts an example decoding result from the high compression, 3-bit encoding example illustrated in Figure 6a implemented in accordance with one embodiment of the invention. The values in Figure 6b are the result of decoding the information placed into the information stream in the example shown in Figure 6a. The decoding table used to determine the values in Figure 6b includes the same quantization table values as the table 400 in Figure 4.

[0054] Figure 7 depicts an exemplary computer system 700 including an imaging device 730 having an encoder 740 implemented in accordance with one embodiment of the invention. Computer system 700 typically includes a bus 702 for communicating information, such as instructions and data. The system further includes a processor 704, coupled to the bus 702, for processing information according to programmed instructions, a main memory 706 coupled to the bus 702 for storing information for processor 704, and an information storage device 708 coupled with the bus 702 for storing information. In the case of a desk-top design for computer system 700, the above components are typically located within a chassis (not shown).

[0055] The processor 704 could be an 80960, 80386, 80486, Pentium<sup>®</sup> processor, Pentium<sup>®</sup> processor with MMX<sup>™</sup> technology, Pentium<sup>®</sup> Pro processor or Pentium<sup>®</sup> II processor made by Intel Corp., among others, including processors that are compatible with those listed above; The processor 704 typically includes a plurality of pipelines for parallel and overlapping execution of operations. The main memory 706 could be a random access memory (RAM) to store dynamic information for processor 704, a read-only memory (ROM) to store static information and instructions for processor 704, or a combination of both types of memory .The processor 704 executes instructions that cause the output of the imaging device 730 to be decoded. The processor may either perform calculations in order to decode the output of the imaging device 730 or may use a decode table stored in the main memory 706. The results of the decoding may be displayed on a display device 710 or stored for later display in the information storage device 708 or in the main memory 706.

[0056] In alternative designs for the computer system 700, the information storage device 708 could be any medium for storage of computer readable information. Suitable candidates include a read-only memory (ROM), a hard disk drive, a disk drive with removable media (e.g., a floppy magnetic disk or an optical disk), or a tape drive with removable media (e.g., magnetic tape), synchronous DRAM or a flash memory (i.e., a disk-like storage device implemented with flash

semiconductor memory). A combination of these, or other devices that support reading or writing computer readable media, could be used.

[0057] The input/output devices of the computer system 700 typically comprise the display device 710, an alphanumeric input device 712, a position input device 714, a communications interface 716, and the imaging device 730, each of which is coupled to the bus 702. The imaging device may comprise a digital camera. If the data storage device 708 supports removable media, such as a floppy disk, it may also be considered an input/output device.

[0058] The imaging device 730 includes the encoder 740 which comprises an encoder quantization table select register 742, a level shift maximum register 748, a scan line buffer 744, encoder quantization tables 745 and 747, and a decoder table 749. The table select register 742 is written to by the processor 704 or other system device in order to select which quantization table is to be used for encoding. The tables 745 and 747 may also be written to by the processor or other system device in order to update the values contained in the tables. Although two encoding tables are shown, any number of encoding tables may be used with the invention. The level shift maximum register 748 is used to store a value that determines the maximum number of level shifts allowed per frame. This register is may also be written to by the processor 704 or other system device. The scan line buffer 744 is used to store at least one previously sampled decoded scan line which is to be used as a predictor for a subsequent scan line. The decoder table 749 is used when decoding scan lines for storage in the line buffer 744. Example tables that may be used to implement tables 745, 747, and 749 are shown in Figures 8a, 8b, and 8c.

[0059] It will be clear to one skilled in the art that the invention can operate upon a wide range of programmable computer systems, not just the example computer system 700.

[0060] Figure 8a depicts an example 4-bit encode table with no under-correction. This example table may be used to implement encoder quantization tables 745 and/or 747 shown in Figure 7. The table shown in Figure 8a uses only absolute values of the calculated difference to be

encoded. This example table contains similar information as the table shown in Figure 4 for 4-bit encoding, but since only absolute values are used, the table is smaller and therefore more efficient to implement in hardware. The absolute difference value must be less than the value in the quantizer (limit) column, with the exception of the "Last" column. The "Last" column contains the upper limits (Max) for the Y, C<sub>R</sub> and C<sub>B</sub> planes. Since the "Last" column contains these upper limits, no "less than" test occurs. When an upper limit is reached, the quantizer from the "Last" column is selected. For all of the columns, if the sign of the calculated difference to be encoded is positive, then the appropriate odd numbered quantizer is selected. If the sign of the calculated difference to be encoded is negative, then the appropriate even numbered quantizer is selected.

[0061] Figure 8b shows an example 4-bit encode table that functions in a manner similar to the table depicted in Figure 8A. The table of Figure 8b, however, has 50% under correction applied. Notice that the table values are twice the values contained in the table of Figure 8a. This example table may also be used to implement encoder quantization tables 745 and/or 747 shown in Figure 7. Further, the table of Figure 8a may be used to implement one of the tables 745 and 747 while the table of Figure 8b may be used to implement the other.

[0062] Figure 8c depicts an example 4-bit decode table that may be used in connection with the tables shown in Figures 8a and 8b. The present example decode table may be used to implement the decoder table 749 as shown in Figure 7. The decode table of Figure 8c also uses absolute values. If the quantizer to be decoded is an even number, then the sign of the decoded value will be negative. If the quantizer to be decoded is an odd number, then the sign of the decoded value will be positive. If the quantizer to be decoded is zero, then the decoded value will be zero.

[0063] Although the tables discussed above in connection with Figures 8a, 8b, and 8c are shown as having particular values, formats, and sizes, the tables may be implemented with varying numbers of levels and quantizers, and may also be implemented with different values. In

particular, the tables of Figures 8a, 8b, and 8c may be implemented as 3-bit encode or decode tables, and also may implement varying amounts of under- correction.

[0064] In the foregoing specification the invention has been described with reference to specific exemplary embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specification and drawings are accordingly to be regarded in an illustrative rather than in a restrictive sense.